

WE CLAIM:

1. For a message system for delivering data in the form of messages between message clients, the message system comprising a server cluster containing a group of message manager nodes with means for storing and distributing messages a method for ensuring operation during node failures and network partitioning, the method comprising:
 - Maintaining, in each running message manager node, a cluster state data set comprising information about the nodes present in the cluster
 - Repeatedly evaluating, in each running one of said message manager nodes, a view state data set comprising information about the server nodes able to communicate with said message manager node,
 - Comparing said cluster state data set with said view state data set and evaluating if all message managers are able to communicate with said message manager node,
 - If not all message managers are able to communicate with said message manager node, determining an operational state for point-to-point style messaging out of at least two different states, where
 - a first one of said at least two different states is an unrestricted state for normal operation, the message managers being in said unrestricted state being allowed to dispatch point-to-point messages sent both prior to and since the last change of the view state of the current network, and
 - a second one of said two different states is a restricted state, the message managers being in said restricted state not being allowed to dispatch point-to-point messages sent prior to the last change of the view state of the current network,

and where the operational state is determined in a manner that of any two message manager nodes that are responsible for dispatching the same point-to-point messages not being able to communicate with each other, at most one can be in an unrestricted operational state.

2. The method of claim 1, further comprising, if not all message managers are able to communicate with said message manager node, determining an operational state for publish/subscribe style messaging out of at least two different states, where

a first one of said at least two different states is a non-retaining state for normal operation, the message managers being in said non-retaining state being allowed to delete publish/subscribe messages as soon as it has determined that all eligible subscribers have received copies of those messages, and

a second one of said two different states is a retaining state, the message managers being in said retaining state not being allowed to delete publish/subscribe messages prior to the expiry of the individual messages or until the message manager changes to the non-retaining state and determines that all eligible subscribers have received copies of those messages,

and where the operational state is determined in a manner that of any two message manager nodes that are responsible for dispatching the same publish/subscribe messages not being able to communicate with each other, both must be in a retaining state.

3. The method of claim 2, wherein the server cluster further comprises a group of client manager nodes, the client manager nodes being programmed differently from the message manager nodes, and wherein

each message manager node at all times attributed a partition state out of at least three different partition states (A, C, D),

where a message manager is attributed a first partition state (A) if it is able to communicate with all other message managers responsible for dispatching the same messages, where it is attributed a further partition state (C) if it is not able to communicate with all other message managers responsible for dispatching the same messages but if it is able to communicate with a majority of the client manager nodes, and where it is attributed a still further partition state (D) otherwise

and where the message manager is determined to be in an unrestricted operational state if it is in the first or in the further partition state (A, C) and determined to be in a restricted operational state otherwise, and where it is determined to be in a non-retaining operational state if it is in the first partition state (A) and to be in a retaining operational state otherwise.

4. The method of claim 2, wherein the server cluster further comprises a group of client manager nodes, the client manager nodes being functionally distinct from the message manager nodes, and wherein

each message manager node is at all times attributed a partition state out of at least four different partition states (A, B, C, D),

where a message manager is attributed a first partition state (A) if it is able to communicate with all other message manager nodes of the cluster that are responsible for dispatching the same set of messages,

where it is attributed a second partition state (B) it is either not able to communicate with a further message manager responsible for dispatching the same kind of messages and is able to communicate with a majority of the client

managers or if its partition state prior to the last view state change was the second partition state and it is still able to not able to communicate with any other message manager responsible for dispatching the same messages and it is able to communicate with at least one client manager,

wherein it is attributed a third partition state (C) if it is not in the first partition state (A) or the second partition state (B) and where at least one further message manager node of the cluster that is responsible for dispatching the same set of messages is either able to communicate with it and it is able to communicate with a majority of the client manager nodes

and wherein it is attributed a fourth partition state (D) if it is not in the first, the second or the third partition state

with the additional restriction that it is not attributed the second state (B) or the third state (C) when, prior to the last change in the view state, it was in the fourth state (D)

and where the message manager is determined to be in an unrestricted operational state if it is in the first, the second or the third partition state (A, B, C) and determined to be in a restricted operational state otherwise, and where it is determined to be in a non-retaining operational state if it is in the first partition state (A) and to be in a retaining operational state otherwise.

5. The method of claim 1 wherein the message manager nodes are attributed an operational state for point-to-point messaging at all times when the server system is running.

6. The method of claim 2, wherein the message manager nodes are attributed an operational state for publish/subscribe messaging at all times when the server system is running.
7. For a message system for delivering data in the form of messages between message clients, the message system comprising a server cluster containing a group of message manager nodes with means for storing and distributing messages a method for ensuring operation during node failures and network partitioning, the method comprising:
- Maintaining, in each running message manager node, a cluster state data set comprising information about the nodes present in the cluster
 - Repeatedly evaluating, in each running one of said message manager nodes, a view state data set comprising information about the server nodes able to communicate with said message manager node,
 - Comparing said cluster state data set with said view state data set and evaluating if all message managers are able to communicate with said message manager node,
 - If not all message managers are able to communicate with said message manager node, determining an operational state for publish/subscribe style messaging out of at least two different states, where
 - a first one of said at least two different states is a non-retaining state, the message managers being in said non-retaining state being allowed to delete publish/subscribe messages as soon as it has determined that all eligible subscribers have received copies of those messages, and
 - a second one of said two different states is a retaining state, the message managers being in said retaining state not being allowed delete

publish/subscribe messages prior to the expiry of the individual messages or the message manager changes to the non-retaining state and determines that all eligible subscribers have received copies of those messages,

and where the operational state is determined in a manner that of any two message manager nodes that are responsible for dispatching the same publish/subscribe messages not being able to communicate with each other, both must be in a retaining state.

8. For a message system for delivering data in the form of messages between message clients, the message system comprising a server cluster divided into sub-clusters containing a group of message manager nodes with means for storing and distributing messages a method for ensuring operation during node failures and network partitioning, the method comprising:
 - Maintaining, in each running message manager node a sub-cluster state data set comprising information about the servers present in the cluster
 - Repeatedly evaluating, in each running one of said message manager nodes, a view state data set comprising information about the server nodes able to communicate with said message manager node,
 - Comparing said sub-cluster state data set with said view state data set and evaluating if all message managers of the sub-cluster are able to communicate with said message manager node,
 - If not all message managers of the sub-cluster are able to communicate with said message manager node, determining a point-to-point operation state out of at least two different states, where
 - a first one of said at least two different states is an unrestricted state, the message servers being in said unrestricted state being allowed to dispatch

point-to-point messages sent prior to and since the last change of the view state of the current network, and

a second one of said two different states is a restricted state, the message servers being in said restricted state not being allowed to dispatch point-to-point messages sent prior to the last change of the view state of the current network.

and where the operational state is determined in a manner that of two message manager nodes of the same sub-cluster not being able to communicate with each other, at most one can be in an unrestricted state.

9. The method of claim 8 further comprising, if not all message managers of the sub-cluster are able to communicate with said message manager node, determining an operational state for publish/subscribe style messaging out of at least two different states, where

a first one of said at least two different states is a non-retaining state for normal operation, the message managers being in said non-retaining state being allowed to delete publish/subscribe messages as soon as it has determined that all eligible subscribers have received copies of those messages, and

a second one of said two different states is a retaining state, the message managers being in said retaining state not being allowed to delete publish/subscribe messages prior to the expiry of the individual messages or until the message manager changes to the non-retaining state and determines that all eligible subscribers have received copies of those messages,

and where the operational state is determined in a manner that of any two message manager nodes of the same sub-cluster not being able to communicate with each other, both must be in a retaining state.

10. For a message system for delivering data in the form of messages between message clients, the message system comprising a server cluster containing a group of message manager nodes with means for storing and distributing messages a method for guaranteeing JMS semantics during node failures and network partitioning, the method comprising:
- Maintaining, in each running message manager node, a cluster state data set comprising information about the nodes present in the cluster
 - Repeatedly evaluating, in each running one of said message manager nodes, a view state data set comprising information about the server nodes able to communicate with said message manager node,
 - Comparing said cluster state data set with said view state data set and evaluating if all message managers are able to communicate with said message manager node,
 - If not all message managers are able to communicate with said message manager node, determining an operational state for point-to-point style messaging out of at least two different states, where
 - a first one of said at least two different states is an unrestricted state for normal operation, the message managers being in said unrestricted state being allowed to dispatch point-to-point messages sent both prior to and since the last change of the view state of the current network, and
 - a second one of said two different states is a restricted state, the message managers being in said restricted state not being allowed to dispatch point-to-point messages sent prior to the last change of the view state of the current network,
- and where the operational state is determined in a manner that of any two message manager nodes that are responsible for dispatching the same point-to-

point messages not being able to communicate with each other, at most one can be in an unrestricted operational state.

and further comprising, if not all message managers are able to communicate with said message manager node, determining an operational state for publish/subscribe style messaging out of at least two different states, where

a first one of said at least two different states is a non-retaining state for normal operation, the message managers being in said non-retaining state being allowed to delete publish/subscribe messages as soon as it has determined that all eligible subscribers have received copies of those messages, and

a second one of said two different states is a retaining state, the message managers being in said retaining state not being allowed to delete publish/subscribe messages prior to the expiry of the individual messages or until the message manager changes to the non-retaining state and determines that all eligible subscribers have received copies of those messages,

and where the operational state is determined in a manner that of any two message manager nodes that are responsible for dispatching the same publish/subscribe messages not being able to communicate with each other, both must be in a retaining state.

11. A message system for delivering data in the form of messages between message clients, comprising a server cluster containing a group of message manager nodes with means for storing and distributing messages a method for ensuring operation during node failures and network partitioning, wherein each message manager node comprises means for
 - Maintaining, in each running message manager node, a cluster state data set comprising information about the nodes present in the cluster

- 57 -

- Repeatedly evaluating, in each running one of said message manager nodes, a view state data set comprising information about the server nodes able to communicate with said message manager node,
- Comparing said cluster state data set with said view state data set and evaluating if all message managers are able to communicate with said message manager node,
- If not all message managers are able to communicate with said message manager node, determining an operational state for point-to-point style messaging out of at least two different states, where

a first one of said at least two different states is an unrestricted state for normal operation, the message managers being in said unrestricted state being allowed to dispatch point-to-point messages sent both prior to and since the last change of the view state of the current network, and

a second one of said two different states is a restricted state, the message managers being in said restricted state not being allowed to dispatch point-to-point messages sent prior to the last change of the view state of the current network,

and wherein said means for determining an operational state out of at least two states are configured in a manner that of two message manager nodes that are responsible for dispatching the same point-to-point messages not being able to communicate with each other, at most one can be in an unrestricted operational state.

12. The system of claim 11, further comprising means for, if not all message managers are able to communicate with said message manager node,

determining an operational state for publish/subscribe style messaging out of at least two different states, where

a first one of said at least two different states is an non-retaining state for normal operation, the message managers being in said non-retaining state being allowed to delete publish/subscribe messages as soon as it has determined that all eligible subscribers have received copies of those messages, and

a second one of said two different states is a retaining state, the message managers being in said retaining state not being allowed to delete publish/subscribe messages prior to the expiry of the individual messages or until the message manager changes to the non-retaining state and determines that all eligible subscribers have received copies of those messages,

and the means for determining an operational state are programmed in a manner that that of any two message manager nodes that are responsible for dispatching the same publish/subscribe messages not being able to communicate with each other, both must be in a retaining state.

13. A computer program product comprising a computer usable medium having computer readable program code means embodied therein for enabling a computer to serve as a message manager in a server cluster, the program product comprising computer readable code means for enabling the computer
 - To maintain, in each running message manager node, a cluster state data set comprising information about the nodes present in the cluster
 - To Repeatedly evaluate, in each running one of said message manager nodes, a view state data set comprising information about the server nodes able to communicate with said message manager node,

- 59 -

- To compare said cluster state data set with said view state data set and to evaluate if all message managers are able to communicate with said message manager node,
- If not all message managers are able to communicate with said message manager node, to determine an operational state for point-to-point style messaging out of at least two different states, where

a first one of said at least two different states is an unrestricted state for normal operation, the message managers being in said unrestricted state being allowed to dispatch point-to-point messages sent both prior to and since the last change of the view state of the current network, and

a second one of said two different states is a restricted state, the message managers being in said restricted state not being allowed to dispatch point-to-point messages sent prior to the last change of the view state of the current network,

and where the means for determining an operational state are programmed in a manner that of any two message manager nodes that are responsible for dispatching the same point-to-point messages not being able to communicate with each other, at most one can be in an unrestricted operational state.

14. A computer program product according to claim 13, wherein said computer readable code means comprise means employing a library written in the Java language and conforming to the Java Message Service API.